

CLAIMS

1. A switching apparatus for sharing input/output endpoints, the switching apparatus comprising:
 - a first plurality of I/O ports, coupled to a plurality of operating system domains (OSDs) through a load-store fabric, each configured to route transactions between said plurality of OSDs and the switching apparatus;
 - a second I/O port, coupled to a first shared input/output endpoint, wherein said first shared input/output endpoint is configured to request/complete said transactions for each of said plurality of OSDs; and
 - core logic, coupled to said first plurality of I/O ports and said second I/O port, configured to route said transactions between said first plurality of I/O ports and said second I/O port, wherein said core logic designates a corresponding one of said plurality of OSDs according to a variant of a protocol, and wherein said protocol provides for routing of said transactions only for a single OSD.
2. The switching apparatus as recited in claim 1, wherein said variant comprises encapsulating an OS domain header within a transaction layer packet that otherwise comports with said protocol.

3. The switching apparatus as recited in claim 1, wherein said first plurality of I/O ports communicates with said plurality of OSDs according to said protocol.
4. The switching apparatus as recited in claim 1, wherein said protocol comprises PCI Express.
5. The switching apparatus as recited in claim 4, wherein said core logic isolates said transactions over a plurality of PCI Express bus hierarchies according to said each of said plurality of OSDs.
6. The switching apparatus as recited in claim 5, wherein a subset of said transactions are routed in accordance with addressing mechanisms for a particular one of said plurality of PCI Express bus hierarchies, said particular one of said plurality of PCI Express bus hierarchies corresponding to a particular one of said plurality of OSDs.
7. The switching apparatus as recited in claim 2, wherein said core logic is configured to associate each of said transactions with a corresponding one of said plurality of OSDs, said corresponding one of said plurality of OSDs corresponding to one or more root complexes.
8. The switching apparatus as recited in claim 7, wherein said transaction layer packet is routed between said second I/O port and said first shared input/output endpoint.

9. The switching apparatus as recited in claim 8, wherein said first shared I/O endpoint is configured to detect said OS domain header and to perform a specified operation according to said protocol exclusively for said one of said plurality of OSDs.
10. The switching apparatus as recited in claim 1, wherein said first shared input/output endpoint designates one of said transactions for a particular one of said plurality of OSDs by encapsulating an OS domain header within a transaction layer packet, and wherein said transaction layer packet is routed to said core logic via said second I/O port.
11. The switching apparatus as recited in claim 10, wherein said core logic is configured to detect and decapsulate said OS domain header from said transaction layer packet, and is configured to route said one of said transactions to said particular one of said plurality of OSDs according to said protocol.
12. A shared input/output (I/O) switching mechanism, comprising:

core logic, configured to enable operating system domains to share one or more I/O endpoints over a load-store fabric, said core logic comprising:

global routing logic, configured to route first transactions to/from said operating system domains, and for routing second transactions to/from said one or more I/O endpoints, wherein each of said second transactions designates an associated one of said operating system domains for which an operation specified by each of said first transactions be performed, and wherein said associated one of said operating system domains is designated according to a variant of a protocol, said protocol providing exclusively a single operating system domain within said load-store fabric.

13. The shared I/O switching mechanism as recited in claim 12, wherein said variant comprises encapsulating an OS domain header within a transaction layer packet that otherwise comports with said protocol.
14. The shared I/O switching mechanism as recited in claim 12, wherein said protocol comprises PCI Express, and wherein said first transactions comport with said protocol.
15. The shared I/O switching mechanism as recited in claim 14, wherein said core logic isolates said first and second transactions over a plurality of PCI Express bus hierarchies according to each of said operating system domains.

16. The shared I/O switching mechanism as recited in claim 15, wherein said first and second transactions are routed in accordance with addressing mechanisms for a particular one of said plurality of PCI Express bus hierarchies, said particular one of said plurality of PCI Express bus hierarchies corresponding to said associated one of said operating system domains.
17. The shared I/O switching mechanism as recited in claim 12, wherein said core logic associates each of said operating system domains with each of a corresponding root complex.
18. The shared I/O switching mechanism as recited in claim 13, wherein said transaction layer packet is routed between said the shared I/O switching mechanism and one of said one or more I/O endpoints.
19. The shared I/O switching mechanism as recited in claim 18, wherein said one of said one or more I/O endpoints is configured to detect said OS domain header and to perform said operation according to said protocol exclusively for said one of said operating system domains.

20. The shared I/O switching mechanism as recited in claim 12, wherein one of said one or more I/O endpoints designates one of said second transactions for a particular one of said operating system domains by encapsulating an OS domain header within a transaction layer packet, and wherein said transaction layer packet is routed to the shared I/O switching mechanism.
21. The shared I/O switching mechanism as recited in claim 20, wherein VMAC logic within the shared I/O switching mechanism is configured to detect and decapsulate said OS domain header from said transaction layer packet, and wherein said core logic routes a corresponding one of said first transactions to said particular one of said plurality of operating system domains according to said protocol.
22. A method for interconnecting independent operating system domains to a shared I/O endpoint within a load-store fabric, comprising:

via first ports, first communicating with each of the independent operating system domains according to a protocol that provides exclusively for a single operating system domain within the load-store fabric;

via a second port, second communicating with the shared I/O endpoint according to a variant of the protocol to enable the shared I/O endpoint to associate a prescribed operation with a corresponding one of the independent operating system domains, said second communicating comprising:

encapsulating an OS domain header within a transaction layer packet that otherwise comports with the protocol, wherein the value of the OS domain header designates the corresponding one of the operating system domains; and

via core logic within a switching apparatus, mapping the independent operating system domains to the shared I/O endpoint.

23. The method as recited in claim 22, further comprising:

over the load-store fabric, first interconnecting the independent operating system domains to the switching apparatus via the first ports; and

over the load-store fabric, second interconnecting the shared I/O endpoint to the switching apparatus via the second port.

24. The method as recited in claim 22, wherein said first communicating comprises:

associating each of the operating system domains with a corresponding root complex.

25. The method as recited in claim 22, wherein said first communicating comprises:

employing PCI Express as the protocol.

26. The method as recited in claim 25, wherein said mapping comprises:

isolating transactions over a plurality of PCI Express bus hierarchies according to each of the operating system domains.

27. The method as recited in claim 26, wherein said mapping further comprises:

routing the transactions in accordance with addressing mechanisms for a particular one of the plurality of PCI Express bus hierarchies, the particular one of said plurality of PCI Express bus hierarchies corresponding to the corresponding one of the operating system domains.

28. The method as recited in claim 22, wherein said second communicating further comprises:

routing the transaction layer packet between the second port and the shared I/O endpoint.

29. The method as recited in claim 28, wherein said second communicating further comprises:

within the I/O endpoint, detecting the OS domain header; and

performing the prescribed operation according to the protocol, wherein said performing is accomplished exclusively for the corresponding one of the operating system domains.

30. The method as recited in claim 22, wherein said second communicating comprises:

within the shared I/O endpoint, designating a transaction for a particular one of the operating system domains by embedding an OS domain header within a transaction layer packet; and

first transmitting the transaction layer packet to the second port.

31. The method as recited in claim 30, wherein said second communicating comprises:

within the switching apparatus, detecting and removing the OS domain header from the transaction layer packet.

32. The method as recited in claim 31, wherein said first communicating comprises:

second transmitting the transaction layer packet to the particular one of the operating system domains according to the protocol.